Road to safe autonomy with data and formal reasoning

Chuchu Fan, Bolun Qi, and Sayan Mitra
{cfan10,bolunqi2,mitras}@illinois.edu
Department of Electrical and Computer Engineering,
University of Illinois at Urbana-Champaign.

arXiv:1704.06406v1 [cs.SY] 21 Apr 2017

We present an overview of recently developed data-driven tools for safety analysis of autonomous vehicles and advanced driver assist systems. The core algorithms combine model-based, hybrid system reachability analysis with sensitivity analysis of components with unknown or inaccessible models. We illustrate the applicability of this approach with a new case study of emergency braking systems in scenarios with two or three vehicles. This problem is representative of the most common type of rear-end crashes, which is relevant for safety analysis of automatic emergency braking (AEB) and forward collision avoidance systems. We show that our verification tool can effectively prove the safety of certain scenarios (specified by several parameters like braking profiles, initial velocities, uncertainties in position and reaction times), and also compute the severity of accidents for unsafe scenarios. Through hundreds of verification experiments, we quantified the safety envelope of the system across relevant parameters. These results show that the approach is promising for design, debugging and certification. We also show how the reachability analysis can be combined with statistical information about the parameters, to assess the risk-level of the control system, which in turn is essential, for example, for determining Automotive Safety Integrity Levels (ASIL) for the ISO26262 standard.

***Index Terms*—Safety verification, Autonomous driving system, ASILs, Risk analysis**

I. Introduction

There is now a race to create commercial autonomous vehicles and more advanced driving assist systems (ADAS). These safety-critical, cyber-physical systems (CPS) use hierarchical control, long supply chains, and increasingly, machine learning algorithms. Existing design and test methodologies are inadequate for providing the needed level of safety assurances. Koopman [1] argues how naïve test-driving for reasonable catastrophic failure rates for a fleet of vehicles can grow to hundreds of billions of miles. At the time of writing, driverless tests from Waymo and Tesla range around 100 million miles and are punctuated by *disengagements*[1]. Precisely measuring risks of these new technologies remains problematic, and the regulations needed for mitigating the risks are indefinite [2].

Could formal verification algorithms provide answers to these challenges? Software model checking, for example, can find design bugs and provide rigorous safety guarantees and have proven to be practical in several domains. The computational problems related to automatically checking the safety of CPS are notoriously difficult (undecidable). Even approximate solutions exist only for relatively simple linear models. Another fundamental problem is that traditional verification methods rely on closed-form, mathematical models of the system (e.g., differential equations and automata). In contrast, automotive systems with hundreds of modules, model-based controllers, machine learning-based units, and fine-tuned lookup tables look less like a model. Could verification algorithms work for systems with possibly incomplete models?

In this article, we report on recent research on data-driven verification that answers this question in the affirmative and we illustrate the promise of this approach with a detailed case study. The key idea is to combine traditional verification of known parts of the model, with *sensitivity-analysis* of the unknown parts that are only executable. Sensitivity analysis gives bounds on how much the states or outputs of a module change, with small changes in the input parameters. In a sequence of papers, we have presented sensitivity analysis algorithms for systems that use available knowledge of models and execution data. Based on these algorithms, we have developed data-driven verification tools C2E2 [3] and DryVR [4]. These tools can certify the safety of an automotive control system under a set of scenarios, or automatically find unsafe scenarios (counter-examples) that violate the safety requirements. In Section II, we present an overview of the inner workings of these techniques. Previous case studies have established the effectiveness of these algorithms in verification of a benchmark engine control system, a NASA-developed collision alerting protocol, and satellite controllers. More recently, DryVR has been used to analyze a range of autonomous and ADAS-based maneuvers [4].

We present a new detailed case study on emergency braking systems in Section II-A. More than 25% of all reported accidents are rear-end crashes [5], of which around 85% happen on straight roads. Emergency braking and forward collision warning systems are becoming standard ADAS features. However, testing safety of such systems can be nontrivial [6]. The safe braking profiles for a sequence of cars on the highway depend on several factors—initial separation, velocities, vehicle dynamics, reaction times, road surface etc. Our data-driven verification tool works with black-box or unknown vehicle models, using which we can prove, for example, that a given braking profile is safe for a set of scenarios characterized by the ranges of initial separation (d) and reaction times (r). For unsafe scenarios, the tool can compute the worst case relative velocity of the collision, which determines the severity of the accident. This type of analysis can be used as a design tool for tuning the braking profiles for different highway speeds, road

[1]A disengagement is an event where the human driver has to take over control from automation to prevent a hazard.

conditions, etc. We analyzed hundreds of scenarios to generate a safety surface that can aid such design and analysis. Finally, we show how data-driven verification can be used for risk analysis. ISO26262 [7] classifies different control subsystems to risk levels (Automotive Safety Integrity Levels or ASILs) and prescribes process-based requirements to reduce those risks to acceptable levels. The risk here is broadly defined as $severity\ of\ accident\ \times\ probability\ of\ occurrence$. Assessing these quantities for complex control systems, however, remains more of an art. In [6] the authors propose a method based on extensive simulations. In contrast, verification gives a provable bound on the severity of accidents for each range of d and r values. We show that this can be combined with statistical information about the distributions of d and r to obtain the risk associated with the system for a given speed and braking profile.

In summary, we present an argument for data-driven formal verification as a foundation for building design automation tools for safe autonomous vehicles and ADAS systems. Specifically, our algorithms combining hybrid system verification and automated sensitivity analysis of black-box models, show promise for both online monitoring and practical design-time and analysis.

II. Verification from Models and Data

Correctness of verification ultimately relies on the underlying system model which may or may not be completely known. We begin by considering dynamical systems—a simple yet very powerful modeling formalism. W have generalized it to more expressive formalisms like switched and hybrid systems, and we refer the interested readers to the article [3].

A *dynamical system* is represented by an ordinary differential equation (ODE):

$$\dot{x}(t) = f(x(t), u(t)) \tag{1}$$

which describes the time-derivative, and hence, the time evolution of a vector x of real-valued state variables (e.g., velocity, torque, steering angles, fuel-flow rates, etc.) with an input signal $u(t)$. Let us fix an input, and denote by $\xi(x_0, t)$ the *solution* of (1) from a particular initial state x_0 at time $t \geq 0$. The exact state is usually hard to estimate; let K be the set of possible initial states, $T > 0$ be the time horizon of interest, and U be a subset of unsafe states. U are the states where speed limits are violated, emissions are excessive, collisions occur, or other bad things happen. The verification question can be written as:

$$\{x \mid \exists x_0 \in K,\ t \in [0, T], \xi(x_0, t) = x\} \ \cap\ U = \emptyset\ ? \tag{2}$$

That is, is there a behaviors of the system from K that enter a bad state within T time. The set on the left hand side is called the *reachset*.

If f is a nonlinear function or is unknown, then computing reachsets can be notoriously difficult. In fact, even if f is known, $\xi(x_0, t)$ may not be computable as a closed form function of x_0 and t. However, it is usually possible to execute or numerically simulate (1), and generate data for $\xi(x_0, 0), \xi(x_0, \tau), \xi(x_0, 2\tau), \ldots, \xi(x_0, T)$ [8]. Data-driven verification algorithms approximate reachsets using this simulation data and sensitivity of $\xi(x_0, t)$ to the changes in x_0.

A. Safety Analysis of Automatic Emergency Braking Systems

An automotive control system model typically consists of several *modes*—for example, cruising, braking, shifting, merging, etc., and the software controller switches between these modes based on sensors and drivers' inputs. This gives rise to a *hybrid* system that combines ODEs with an automaton that defines the allowed mode switches.

Consider for example an Automatic Emergency Braking (AEB) system (Figure 1): $\mathsf{Car1}, \mathsf{Car2}$ and $\mathsf{Car3}$ are cruising down the highway with zero relative velocity and certain initial relative separation; $\mathsf{Car1}$ suddenly switches to a braking mode and starts slowing down according to a certain deceleration profile. Irrespective of whether $\mathsf{Car2}$ is human-driven, AEB-equipped, or fully autonomous, certain amount of time elapses, before $\mathsf{Car2}$ switches to a braking mode. We call this the *reaction time* r. Similarly, the mode switch for $\mathsf{Car3}$ happens after a delay. Obviously, $\mathsf{Car2}$'s safety is in jeopardy: if it brakes "too hard" it will be rear-ended and if it is "too gentle" then it would have a forward collision. The envelope of safe (no collision) behaviors depend on all the parameters: initial separations, velocities, braking profiles, and reaction times. It is easy to see that if we can solve the safety verification problem described above, then we can also compute this envelope and determine whether a given AEB system is safe over a range of scenarios.

Delving deeper into this system, we remark that the exact vehicle dynamics and braking profiles are typically complex or unknown, but we can simulate the system. For the results in Sections IV,V, for example, we used a black-box vehicle model from Mathworks® demo [9]. Each car can be viewed a hybrid system: it has several continuous state variables: deceleration rate $a(t)$, velocity $v(t)$ and position $s(t)$ and two modes: $cruise$ and $brake$. In the $cruise$ mode, the car maintains a constant speed, and in the $brake$ mode, it decelerates according to a certain braking profile for $a(t)$, which is an input to the system. The initial set K of the system is defined by the uncertainties in the initial vehicle velocities $(v_1(0), v_2(0), v_3(0))$ and the initial separations d_{12}, d_{23} between the vehicles. The unsafe set U corresponds to states where there is a collision, that is, the separation between a pair of cars is less than some threshold. In case of collisions, we would be interested in the maximum possible relative velocity just before the collision, which strongly influences the severity of the accident.

The key parameters we will consider in this paper are the reaction time or the delay r in switching to the braking mode, the initial separation d between the cars.

III. Sensitivity Analysis and Verification Algorithms

Data-driven verification algorithms rely on computing reachsets from models and simulation data. First we discuss a situation where a lot of information about the system model is available, and this is exploited by the verification algorithm.

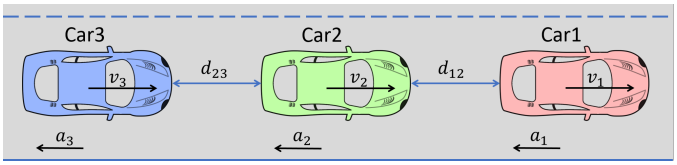

Fig. 1: Cars cruising and braking in a single lane configuration.

Then step by step we will drop these assumptions and arrive at algorithms that only use the system as a black-box.

A. Verification Algorithm

As mentioned in the introduction, the key idea is to use simulations to determine the sensitivity of the system's solution $\xi(x_0, t)$ to changes in initial conditions x_0. Formally, sensitivity is quantified by the the following notion discrepancy.

Definition 1. *A uniformly continuous nonnegative function β is a* discrepancy function *of (1) if (a) for any pair of states x, x', and any time $t > 0$,*

$$\|\xi(x,t) - \xi(x',t)\| \leq \beta(x, x', t), \text{and} \tag{3}$$

(b) for any t, as $x \to x'$, $\beta(\cdot, \cdot, t) \to 0$.

Recall, the safety verification problem of Equation (2). Assuming that a discrepancy function β is available for the system (1), the safety verification algorithm for (1) proceeds as follows:

1) Compute a δ-cover $C = \{x_i\}_{i=1}^{k}$ of the initial set K, i.e., $K \subseteq \cup_i B_\delta(x_i)$.
2) For each $x_i \in C$, a simulation $\psi(x_i)$ from x_i is computed.
3) The simulation $\psi(x_i)$ from x_i is expanded by a factor given by the discrepancy, such that this expanded simulation $R(\psi(x_0), \beta, \delta)$ over-approximates all the states reachable from a δ-neighborhood of x_i over the time horizon T.
4) If this reach set $R(\psi(x_0), \beta)$ is disjoint from the unsafe set U then x_i is removed from the cover C. Else if any interval of the simulation $\psi(x_0)$ is contained in U then output **Unsafe** and $\psi(x_0)$ serves as a counter-example or bug trace. Otherwise, if neither case holds, then x_i is replaced in C by a finer cover of the δ-neighborhood of x_i.
5) If the cover C becomes empty, then output **Safe**.

This simple algorithm computes increasingly finer covers of K until the reachsets $R(\psi(x_0), \beta, \delta)$ from each of the elements in the cover are inferred to be disjoint from the unsafe set U or a counter-example simulation is discovered. The second property of the discrepancy function ensures that as the elements in the cover become finer, the over-approximation of the computed reachsets becomes more precise. A simple 2-D illustration of the verification algorithm in action is shown in Figure 2.

Essentially the same idea can be made to work for switched and hybrid models like the emergency braking scenarios described in Section III-B (see [3] for details). The main complication is that because of the over-approximations in the computed reachsets, we have to keep track of spurious mode changes.

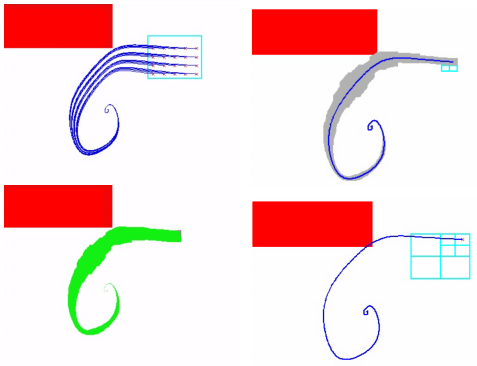

Fig. 2: Conceptual demonstration of data-driven verification. Red rectangle: Unsafe states, Cyan rectangle: Uncertainty space. Simulations (blue lines) cannot guarantee safety, but simulation data together with sensitivity analysis can give coverage (grey region) to prove safety (green region) or help identify bugs quickly.

The algorithm has been proved to be sound and relatively complete in [3]. Soundness means that whenever the Algorithm returns **Safe** or *Unsafe*, then the system is indeed safe or unsafe, respectively. For example, if the algorithm returns safe for a AEB scenario with certain braking profiles, then there is guaranteed to be no collisions until all the cars stop. If the algorithm returns **Unsafe**, then there are particular parameter values, such as the initial separations and reaction times, that leads to a collision. In this case, our tool will compute an upper-bound on the maximum relative velocity of all collisions that can arise from the considered range of parameter values. Relative completeness means the algorithm will always terminate whenever the system is either robustly safe or robustly unsafe.

B. Computing Discrepancy

The key ingredient for the above verification algorithm is a good discrepancy funciton β. One can get an exponentially growing discrepancy function $\beta(x, x', t) = ||x - x'||e^{Lt}$, where L is a Lipshitz constant for f, but even for moderately large L, the reachset over-approximations $R(\psi(x_0), \beta, \delta)$ blow-up, and the verification algorithm would get clogged by large number of refinements.

Methods for computing tight discrepancy functions for linear ODEs were presented in [3], [10], but the problem remained open for general nonlinear models. In [11], an approach was presented for nonlinear systems that aimed to compute local discrepancy functions that are relevant only over small parts of the state space. It was shown that this preserves the soundness and relative completeness of the verification algorithm, and this approach lies at the core of our first data-driven verification tool for C2E2 [3].

Previous methods for computing discrepancy in [10], [11] rely on availability of a closed-form system model (i.e. the dynamical function f). In many practical control systems, the model is at least partly unavailable or it is too complicated for deriving a closed-form description. In this case, hybrid

control systems can be described by combining a black-box simulator for trajectories and a white-box transition graph specifying mode switches. When we only have access to a black-box simulator, a probabilistic algorithm can be used to learn the parameters of exponential discrepancy functions from simulation data. This is the basis for our new data-driven verification tool DryVR [4].

The DryVR algorithm transforms the problem of learning the parameters of the discrepancy function to the problem of learning a linear separator for a set of points in 2-dimensions that are obtained from transforming the simulation data. The idea of the algorithm works as follows: (1) Draw m samples of initial states $x_i, i = 1, \ldots, m$ from initial set K, (2) Simulate the black-box simulator to get sampled traces $\xi(x_i, t_k), i = 1, \ldots, m, k = 0, \ldots, N$, where $t_N = T$ is the time bound, (3) Minimize $\int_{t=0}^{T} ce^{\gamma t}$ such that $ce^{\gamma t}$ with c, γ a scalar value is a valid discrepancy function for any pair of traces $\xi(x_i, t), \xi(x_j, t), i, j = 1, \ldots, m$ and for any time $t = t_k, k = 0, \ldots, N$. With the PAC-learnability of concepts with low VC-dimension, it can be shown that for $\delta, \epsilon > 0$, if $m \geq \frac{1}{\epsilon} \ln \frac{1}{\delta}$, then with probability $\leq \delta$, the constructed discrepancy function $ce^{\gamma t}$ does not work for more than ϵ fraction of the points in the initial set K. Assuming that the discrepancy function is correct, the DryVR verification algorithm gives the same soundness and relative completeness quarantees as stated in Section III-A. Our experiments suggest that a few dozen simulation traces are adequate for learning discrepancy functions with nearly 100% correctness, for typical automotive models.

C. Determining Severity of Collisions using Reachability

Let us consider the AEB systems discussed in Section II-A and see how reachability analysis can guarantee safety of scenarios and compute worst-case collision velocities.

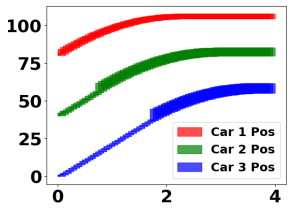

(a) Safe case: reachtubes of position are separated by a distance ≥ 2 for any time.

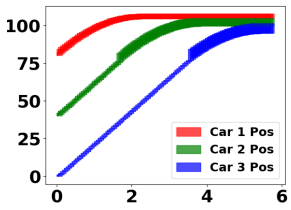

(b) Unafe case: at least a pair of reachtubes of position for some time are contained in the unsafe region (≤ 2 separations).

Fig. 3: Safety of the AEB system. Horizontal axis is time in seconds, vertical axis is position in meters.

Fix the initial velocities and braking profiles for all the cars, fix the ranges for initial separations and reaction times, if DryVR returns safe (and the learned discrepancy is correct), then it means the distance between any two cars is always larger than the threshold value $\theta = 2$ meters at all times, before the cars stop. Figure 3a shows the projection of the reachsets plotted against time for the entire range of initial conditions; the range of positions for Car1 (red),Car2 (green) and Car3(blue) are separated by at least 2 meters. If DryVR returns unsafe, then it also computes parameter values (initial separations d_{12}, d_{23}, reaction times r_2, r_3) that lead to a state where the cars have less than 2 meters separation. In Figure 3b, the reachsets for position overlap indicating a collision and in this case the tool also over-approximates the worst case relative velocity in the collision. For example, in the particular example the worst case collision velocity between Car1 and Car2 is $9.0(m/s)$.

IV. Risk analysis for ASIL

Reachability analysis can be used for determining risk levels of an automotive control system. According to ISO 26262 ASIL classification, risk is broadly defined as *severity of accident* $\times$ *probability of occurence*. For the AEB system with 2 cars, the severity is largely determined by the relative velocity of collision, which is approximated from the above reachability analysis.

The probability of occurrence depends on the probability distributions on the parameters (d, r, etc.). In general, these distributions can be complicated. As a starting point, the preliminary study presented in [6], use empirical observations to construct distributions on initial separation (d) which turns out to be a skewed Gaussian with the mean dependent on the car speed. Similarly, the reaction time distribution is also a skewed Gaussian. Examples of such distributions are built using [6], [12] and shown in Figures 4a and 4b, where the separation d ranges over $[40, 50]$ meters and reaction time r ranges over $[0.7, 2.4]$ seconds.

We analyze the risk by dividing $[40, 50]$ in to 10 consecutive small intervals $[d_l^i, d_u^i], i = 1, \ldots, 10$, and $[0.7, 2.4]$ into 17 consecutive intervals $[r_l^i, r_u^i], j = 1, \ldots, 17$. For each region consists of small intervals of d and r, we use DryVR to verify safety or compute the worst case collision velocity. To compute the probability of the accident occurring, we need to compute the probability that each parameter lies in the given range. For distributions shown in Figures 4a and 4b, the probability of the region $d \in [d_l^i, d_u^i], r \in [r_l^j, r_u^j]$ is $Pr(d \in [d_l^i, d_u^i]) \times Pr(r \in [r_l^j, r_u^j])$ if we assume the events $d \in [d_l^i, d_u^i])$ and $r \in [r_l^j, r_u^j]$ are independent. For example, $Pr(41 \leq d \leq 42, 1.0 \leq r \leq 1.1) = 0.19 \times 0.139 = 0.026$.

With the given braking profile and initial velocity of both cars, we can compute the worst case relative velocity for region of d and r. We report the results in Figure 4c. The numbers correspond to each rectangle in the figure are the worst case relative velocities. For example, for the case $d \in [40, 41]$ and $r \in [2.3, 2.4]$, the worst case relative velocity v_c is $17.5(m/s)$. We also plot the heat map of risks as the background of Figure 4c, where the green rectangles with number 0 correspond to the safe cases. Combined with the probability of occurrence, we can compute the expected velocity in the collision for Figure 4 to be $E[v_c] = \sum_{i=1}^{10} \sum_{j=1}^{17} Pr(d \in [d_l^i, d_u^i]) \times Pr(r \in [r_l^j, r_u^j]) \times v_c(i, j) = 2.86(m/s)$. Therefore, for AEB system with given braking profile and initial veclocity for each car, and given distributions for initial separations and reaction times, we can compute the risk defined in ASIL as the expected worst case relative velocity for the collisions.

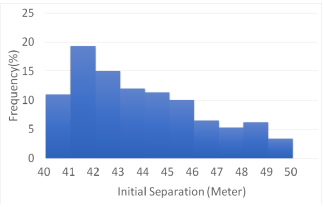

(a) Initial separation distribution

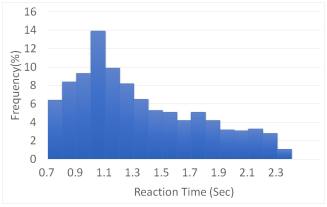

(b) Reaction time distribution

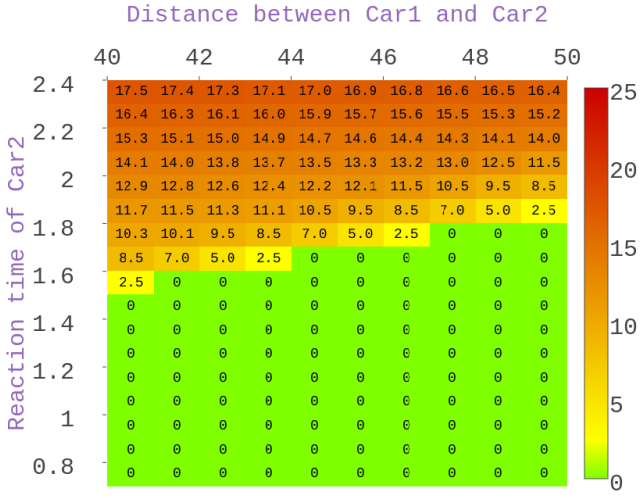

(c) Worst case relative velocities (m/s) for the collisions. Braking profiles are fixed (Car1: mild brake, Car2: medium brake) and initial velocities are $30(m/s)$.

Fig. 4: AEB of two cars: probability and severity

V. Integrated Safety Analysis

Data-driven verification can be used to gain detailed insights about the safety of autonomous and ADAS systems under different scenarios and parameter variations. For the emergency braking system with two and three vehicles, we have analyzed hundreds of experiments; the summary of the worst-case relative collision velocities computed from these experiments are shown in Figures 5.

We consider 3 different braking profiles for each vehicle: mild, medium and hard. The average deceleration rate increases from mild to hard. The risk analysis can also be applied to any braking profile like the threshold braking and cadence braking used in Anti-lock Braking Systems. Figures 5a and 5d show the collision heat maps with fixed initial velocities but changing braking profiles for two cars. From Figures 5a, 5b and 5d, we observe that if the lead and the following cars have the similar level of braking, the safe regions are nearly invariant. However, with the increasing of the braking level, the severity (relative velocities) of collisions also increase. Comparing with Figure 5d, we can see that if the lead car brakes harder than the follower, then as expected, the safety regions shrink rapidly. Moreover, the collisions are more severe than those in the previous case with both cars braking equally hard. If the lead car brakes more gently (Figure 4c), then the severity reduces quickly. Therefore, qualitatively, it is safer for the following car to choose a braking profile harder or equal to the braking profile of the lead car.

Figures 5d-5f show a sequence of collision heat map with fixed braking but changing initial velocities. As expected, both the area of the unsafe regions and severity of collisions decrease with reduction of the initial velocities. The analysis enables us to prove that, for example, the system is safe when the initial velocities of both cars are less than $17(m/s)$ for the given braking profiles and reaction times.

For the system with three cars, we consider scenarios with 4 parameters: the initial separations d_{12}, d_{23} and reaction times r_2, r_3. For visualizing the risk, we fix the range of 2 parameters while varying the others. Fixing the reaction times of both Car2 and Car3 to be within the range $[1.8, 1.9](s)$, we analyze the change of safety envelope with respect to the change of d_{12} and d_{23}. Figure 5g shows that the system is collision-free only when both the distances d_{12} and d_{23} are large enough. Compare Figure 5g with Figure 5e when all the cars have the same initial velocities and braking profiles. We can see the when the reaction time is between $[1.8, 1.9](s)$, the safe distance change from $d > 44(m)$ for system with two cars to $d_{12} > 47(m), d_{23} > 49(m)$ for system with three cars. Therefore, with the increase of number of cars in a chain, the "safe" distance between any pair of cars increases as well.

Next, fixing the distance d_{12}, d_{23} to be within the range $[44, 45](m)$, we analyze the change of safety envelope with respect to the change of reaction time r_2, r_3. Figure 5h shows that the cars are collision free only if both Car2 and Car3's reaction time are short enough. Compared with Figure 5e again, when the distance between the cars are between $[44, 45](m)$, the safe reaction time change from $r < 1.9(s)$ for the two cars scenario to $r_2 < 1.7(s), r_3 < 1.6(s)$ for three cars scenario. Both Figure 5g and 5h show quantitatively that the safety envelope shrinks with the increase of number of cars in the system.

As the running time for each scenario is $3 - 5$ seconds on a standard laptop, it takes $5 - 30$ minutes to generate a heat map, which suggest that similar analysis could be applied to more complicated scenarios with larger number of modes, parameters, and more sophisticated ADAS systems.

VI. Conclusion

We sketched recent developments in verification tools of CPS that combine formal reasoning with simulation data to effectively prove safety or estimate worst case accidents for automotive control systems. Our case study with emergency braking show that designers can use the tool for analyzing autonomous driving and ADAS features under a variety of traffic scenarios. Engineering the tools to scale to bigger scenarios with many more modes and vehicles while achieving near real-time performance will be an natural and important next step.

Acknowledgement

We thank Joe Miller for pointing us to the emergency braking case study. This case study would be impossible without our collaboration with Mahesh Viswanathan on DryVR.

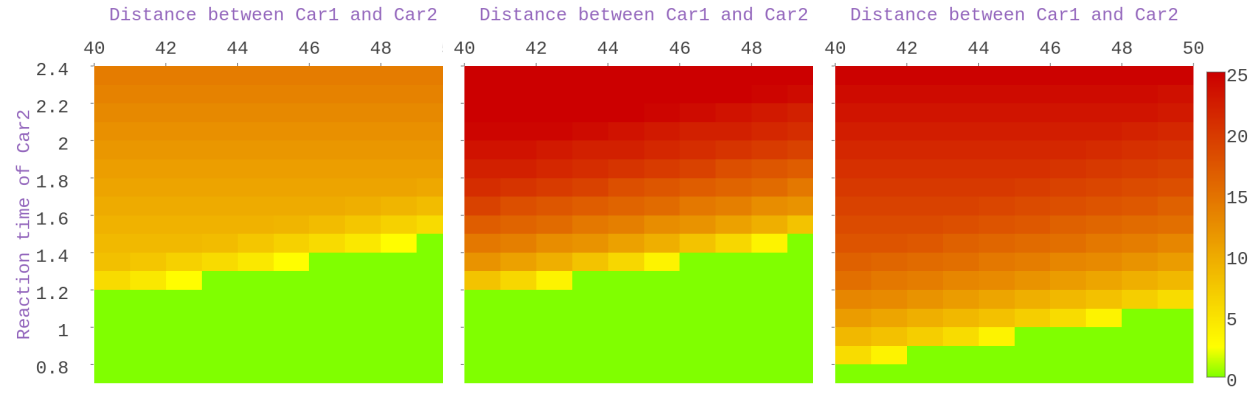

(a) Car1 and Car2: mild brake, initial velocity: $30(m/s)$

(b) Car1 and Car2: hard brake, initial velocity: $30(m/s)$

(c) Car1: medium brake, Car2: mild brake, initial velocity: $30(m/s)$

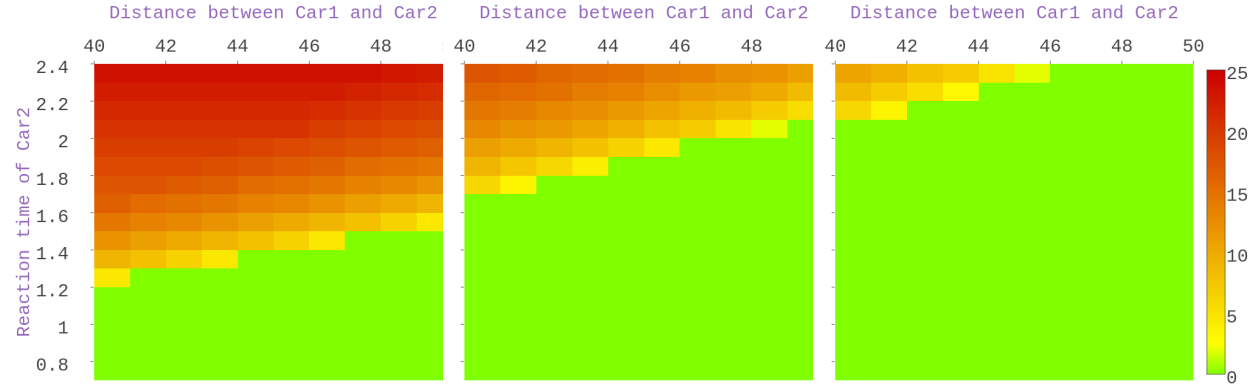

(d) Car1 and Car2: medium brake, initial velocity: $30(m/s)$

(e) Car1 and Car2: medium brake, initial velocity: $22(m/s)$

(f) Car1 and Car2: medium brake, initial velocity: $18(m/s)$

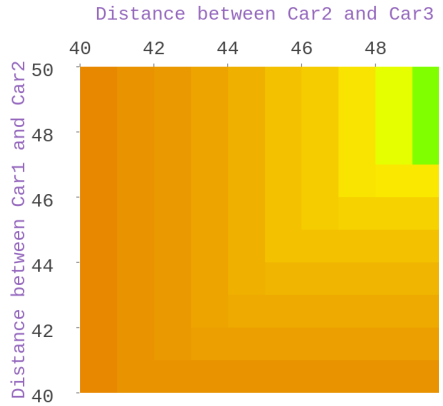

(g) Initial separation: d_{23} vs d_{12}. Fix the reaction time of both Car2 and Car3: $r_2, r_3 \in [1.8, 1.9](s)$.

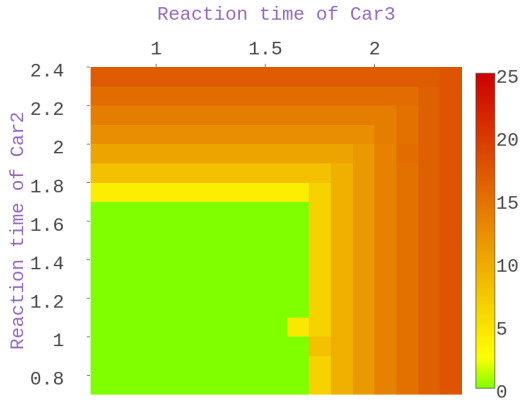

(h) Reaction time: r_3 vs r_2. Fix the initial separation of both pairs of cars: $d_{12}, d_{23} \in [44, 45](m)$

Fig. 5: Top row: two cars with different braking profiles and fixed initial velocities. Middle row left to right: two cars with decreasing velocities and fixed braking profiles. Bottom row: three cars with each car's deceleration is medium hard, and initial velocity is $22(m/s)$.

REFERENCES

[1] P. Koopman and M. Wagner, "Challenges in autonomous vehicle testing and validation," *SAE International Journal of Transportation Safety*, vol. 4, no. 2016-01-0128, pp. 15–24, 2016.

[2] National Highway Traffic Safety Administration and others, "Preliminary statement of policy concerning automated vehicles," *Washington, DC*, pp. 1–14, 2013.

[3] P. S. Duggirala, S. Mitra, M. Viswanathan, and M. Potok, "C2E2: A verification tool for stateflow models," in *Tools and Algorithms for the Construction and Analysis of Systems*, pp. 68–82, Springer, 2015.

[4] C. Fan, B. Qi, S. Mitra, and M. Viswanathan, "Dryvr: Data-driven verification and compositional reasoning for automotive systems," *CoRR*, vol. abs/1702.06902, 2017.

[5] K. Kodaka, M. Otabe, Y. Urai, and H. Koike, "Rear-end collision velocity reduction system," tech. rep., SAE Technical Paper, 2003.

[6] S. Fabris, "Method for hazard severity assessment for the case of undemanded deceleration," *TRW Automotive, Berlin*, 2012.

[7] "Road vehicles - Functional safety," standard, International Organization for Standardization (ISO), Geneva, Switzerland, Nov. 2011.

[8] N. Nedialkov, "VNODE-LP: Validated solutions for initial value problem for ODEs," tech. rep., McMaster University, 2006.

[9] Mathworks, "Simple 2D Kinematic Vehicle Steering Model and Animation." https://www.mathworks.com/matlabcentral/fileexchange/54852-simple-2d-kinematic-vehicle-steering-model-and-animation?requestedDomain=www.mathworks.com.

[10] A. Donzé, "Breach, a toolbox for verification and parameter synthesis of hybrid systems," in *International Conference on Computer Aided Verification*, pp. 167–170, Springer, 2010.

[11] C. Fan, J. Kapinski, X. Jin, and S. Mitra, "Locally optimal reach set over-approximation for nonlinear systems," in *Proceedings of the 13th International Conference on Embedded Software*, p. 6, ACM, 2016.

[12] J. Piao and M. McDonald, "Low speed car following behaviour from floating vehicle data," in *IEEE IV2003 Intelligent Vehicles Symposium. Proceedings (Cat. No.03TH8683)*, pp. 462–467, June 2003.

Chuchu Fan is a PhD candidate at Electrical and Computer Engineering Department, University of Illinois at Urbana-Champaign. She received her Bachelor degree from Automation Control Department, Tsinghua University in 2013. Her research interests are in verification of nonlinear hybrid systems.

Bolun Qi is a master student at Computer Science Department, University of Illinois at Urbana-Champaign, where he received his Bachelor degree as well in 2016. His research interests are in development of formal verification tools.

Sayan Mitra is an Associate Professor of Electrical and Computer Engineering at the University of Illinois at Urbana-Champaign. His research interests are in formal methods, distributed systems and hybrid control systems with applications in automotive, medical, and robotic systems. He received a PhD from MIT and held a post-doctoral fellowship at California Institute of Technology. He has held visiting positions at Oxford University and the Air Force Research Laboratory at New Mexico. He received the National Science Foundation's CAREER Award, Air Force Office of Scientific Research Young Investigator Award, and the IEEE-HKN C. Holmes MacDonald Outstanding Teaching Award.